TRUNCATED LINEAR MODELS FOR FUNCTIONAL DATA

Peter Hall
University of Melbourne

Giles Hooker
Cornell University

arXiv:1406.7732v1 [stat.ME] 30 Jun 2014

Abstract. A conventional linear model for functional data involves expressing a response variable Y in terms of the explanatory function $X(t)$, via the model: $Y = a + \int_{\mathcal{I}} b(t)\, X(t)\, dt + \text{error}$, where a is a scalar, b is an unknown function and $\mathcal{I} = [0, \alpha]$ is a compact interval. However, in some problems the support of b or X, $\mathcal{I}_1$ say, is a proper and unknown subset of $\mathcal{I}$, and is a quantity of particular practical interest. In this paper, motivated by a real-data example involving particulate emissions, we develop methods for estimating $\mathcal{I}_1$. We give particular emphasis to the case $\mathcal{I}_1 = [0, \theta]$, where $\theta \in (0, \alpha]$, and suggest two methods for estimating a, b and θ jointly; we introduce techniques for selecting tuning parameters; and we explore properties of our methodology using both simulation and the real-data example mentioned above. Additionally, we derive theoretical properties of the methodology, and discuss implications of the theory. Our theoretical arguments give particular emphasis to the problem of identifiability.

Key words and phrases. Consistency, functionally equivalent models, identifiability, mean squared prediction error, regression, slope function, statistical smoothing, support interval.

Short title. Truncated functional regression.

1. INTRODUCTION

1.1. Linear model for functional data. The linear model for functional regression has the form

$$Y = a + \int_{\mathcal{I}} b\, X + \epsilon\,, \tag{1.1}$$

where independent observations of the pairs (X, Y) are made, X is a random function recorded on the interval $\mathcal{I}$, a and Y are scalars, b is a function defined on $\mathcal{I}$, and ϵ denotes an experimental error with zero mean. In the case of a truncated linear model there are practical reasons to believe that Y depends on X only through the values taken by X on a subinterval $\mathcal{I}_1 = [u, v]$, say, of $\mathcal{I}$.

Therefore, in place of (1.1), we ask that

$$Y = a + \int_u^v b\, X + \epsilon\,. \tag{1.2}$$

If u and v in (1.2) are included, along with a and b, among the unknowns in the model at (1.2), then the model is no longer linear. It is, of course, an example of functional linear regression, which we consider immediately below.

1.2. General functional regression. In regression we observe independent replicates of the data pair (X, Y), and the relationship between X and Y is modelled as

$$Y = g(X) + \epsilon\,. \tag{1.3}$$

Here g is a real-valued function, or a functional if X is a function, and the experimental error ϵ satisfies

$$E(\epsilon \,|\, X) = 0\,. \tag{1.4}$$

Suppose we can parametrise g, either in a conventional sense where only a finite number of parameters are involved, or in a nonparametric setting where the number of parameters is countably infinite. In the first of these contexts we often estimate unknown parameters by minimising an empirical version of the mean squared prediction error,

$$D_1(g_{\rm mod}) = E\{Y - g_{\rm mod}(X)\}^2\,, \tag{1.5}$$

where $g_{\rm mod}$ represents a model that, in cases where X is a random function, might be particularly complex. In nonparametric settings we typically do the same, after disregarding all but m, say, of the unknown parameters, and letting m increase with sample size.

The attraction of minimising $D_1(h)$ is that, under a condition such as (1.4), $D_1(g_{\rm mod})$ equals the mean squared difference between the true $g(X)$ and the model $g_{\rm mod}(X)$, the latter expressed as function of unknown parameters, plus the quantity $E(\epsilon^2)$, which does not depend on the model. Of course, this result does not require the full force of (1.4); it needs only the property that $g(X) - g_{\rm mod}(X)$ and ϵ are uncorrelated, which follows from (1.4).

The fact that $g_{\rm mod}$ can be particularly complex motivates consideration of simpler functions, or functionals, alternative to both $g_{\rm mod}$ and g. These alternatives might be far too simple to capture the true g in any detail, but they can be much simpler to analyse, and hence also simpler to use for prediction. Importantly, and as we shall show in section 2, these alternative functions include the truncated linear

model at (1.2). This property leads to simple results about the identifiability of that model; see section 2.2.

1.3. Literature survey. Methodology for the functional linear model was discussed in Chapter 10 of Ramsay and Silverman (2002), and Chapter 12 of Ramsay and Silverman (2005). Cardot et al. (1999) made a particularly early contribution to the field. Cardot et al. (2003), and Zhang and Chen (2007), discussed the impact of smoothing on inference in the functional linear model; Crambes et al. (2008, 2009), and Maronna and Yohai (2013), introduced methods based on smoothing splines; Baíllo (2007) suggested kernel techniques and made comparisons with parametric approaches; James et al. (2009) developed variable selection ideas; Mas and Pumo proposed an alternative formulation of the functional linear model; He et al. (2010) introduced techniques based on canonical analysis; Yuan and Cai (2010) suggested a method founded on reproducing kernel Hilbert space analysis; Ferraty et al. (2012) discussed presmoothing methods; and Comte and Johannes (2012), Johannes and Schenk (2012) and Cai and Zhou (2013) treated methods for adaptive smoothing in functional linear regression. Fan and Zhang (2000), Fang et al. (2005) and Wu et al. (2010), among others, developed methodology for functional linear regression in the context of longitudinal data analysis; Cai and Hall (2006) and Apanasovich and Goldstein (2008) addressed mean squared prediction error in functional linear regression; and Cardot et al. (2007), Hall and Horowitz (2007), Li and Hsing (2007), Cai and Yuan (2012) and Johannes and Schenk (2013) discussed convergence rates of estimators of a and b in (1.1).

1.4. Summary. We begin in section 2 by exploring, in general cases including those where X is a function, the class of all candidates for the regression mean g in (1.3). In particular, in the setting of the truncated linear model at (1.2), we show that the intercept a and slope function b are identifiable in particularly general circumstances. This general perspective underpins our development, in section 3, of methodology for estimating a, b and the support interval $\mathcal{I}_1 = [u, v]$. We suggest two methodologies, introduced in parts A and B, respectively, of section 3.1; in section 3.2 we illustrate the application of those techniques, depending as they do on tuning parameters; and in section 3.3 we introduce methods for choosing the tuning parameters. Sections 4, 5 and 6 illustrate properties of our methodology through simulation analysis, by application to real data, and through theoretical

development, respectively. Technical arguments are deferred to appendix A.

2. GENERAL REGRESSION MODELS

2.1. General regression and correlation. Let the regression mean g be as at (1.3), and let the alternatives to g be members, h say, of a class $\mathcal{H}$. They are appropriate regression models, even when they are incorrect (that is, even when $g \notin \mathcal{H}$), provided that the version of (1.5) when $g_{\rm mod}$ is replaced by h can be written as $E\{g(X) - h(X)\}^2$, plus a quantity that does not depend on h. This property, if it were to hold, would reflect the lack of correlation between $g(X)$ and ϵ discussed in the previous paragraph, and it is captured by the following:

> the version of the experimental error ϵ that is implicit in the form of (1.3) for the new regression problem, is uncorrelated with the fitted mean $h(X)$. (2.1)

We claim that the following constraint is sufficient for (2.1):

> for each $h \in \mathcal{H}$, and all constants $c_1 \in (-\infty, \infty)$ and $c_2 > 0$, the function $c_1 + c_2\, h$ is also in $\mathcal{H}$. (2.2)

Condition (2.2) is equivalent to asking that each h can be rescaled and recentred at will, without leaving $\mathcal{H}$. Of course, we require that the quantities $g(X)$, $h(X)$ and ϵ have finite variance:

$$E\{g(X)^2\} + E(\epsilon^2) < \infty \text{ and, for all } h \in \mathcal{H},\ E\{h(X)^2\} < \infty\,. \tag{2.3}$$

Theorem 2.1. *If* (1.4), (2.2) *and* (2.3) *hold, then any function* $h = h_0 \in \mathcal{H}$ *that minimises*

$$D_2(h) = E\{Y - h(X)\}^2 \tag{2.4}$$

satisfies

$$E\{g(X)\} = E\{h_0(X)\}\,, \quad E\big\{h_0(X)^2\big\} = E\{g(X)\, h_0(X)\}\,. \tag{2.5}$$

To appreciate why (2.1) follows from Theorem 2.1, note that if we choose h_0 to minimise $D_2(h)$ at (2.4), and treat $h_0(X)$ as the new version of $g(X)$, then the model error alters from ϵ to $\epsilon_0 = \epsilon + g(X) - h_0(X)$, and in this setting we can write (1.3) equivalently as $Y = h_0(X) + \epsilon_0$. Property (2.1) asks that the new $g(X)$, i.e. $h_0(X)$, and the new error, i.e. ϵ_0, be uncorrelated, and it follows from (2.5) that this is indeed the case.

Of course, if X is a random function then (2.2) holds if $\mathcal{H}$ represents the functional linear model. This is a major attraction of that model—not only is it relatively simple to analyse, but it remains valid as a regression model since its associated noise is uncorrelated with the signal, even if h is not identical to the more intricate functional g.

2.2. Identifiability of the truncated linear model. Recall the definition, at (1.2), of the truncated linear model for functional data. In the present section we show that, under a mild condition on the distribution of X (see (2.8) below), a truncated linear model can be identified from data. Consider the possibility that there exists an alternative, functionally equivalent linear model, where the intercept and slope function a and b are replaced by a_1 and b_1, respectively, and the interval $[u, v]$ is instead $[u_1, v_1]$:

$$P\biggl(a + \int_u^v b\,X = a_1 + \int_{u_1}^{v_1} b_1\,X\biggr) = 1\,, \tag{2.6}$$

where $u, v, u_1, v_1 \in \mathcal{I}$, $u < v$ and $u_1 < v_1$.

Suppose too that $\int_{\mathcal{I}} E(X^2) < \infty$, and let

$$K(t_1, t_2) = \mathrm{cov}\{X(t_1), X(t_2)\} = \sum_{j=1}^{\infty} \omega_j\,\phi_j(t_1)\,\phi_j(t_2) \tag{2.7}$$

denote the singular-value decomposition of the covariance function K, where $\omega_1 \geq \omega_2 \geq \ldots$ are eigenvalues, and $\phi_1, \phi_2, \ldots$ are the associated eigenfunctions, of the linear operator with kernel K. We assume that:

> the linear operator with kernel K is of full rank in $L_2(\mathcal{I})$, in the sense that each $\omega_j \neq 0$ and the sequence $\phi_1, \phi_2, \ldots$ is complete in the class of square-integrable functions on $\mathcal{I}$. (2.8)

Theorem 2.2. *If* (2.6) *and* (2.8) *hold then* $a = a_1$ *and*

$$b(t)\,I(t \in [u, v]) = b_1(t)\,I(t \in [u_1, v_1]) \tag{2.9}$$

for almost all $t \in \mathcal{I}$.

To appreciate the implications of Theorem 2.2, suppose the function b is strictly positive on (u, v) and vanishes on $\mathcal{I} \setminus (u, v)$, and take u_1 and v_1 to be respectively the supremum and infimum of all candidate values t, for u_1 and v_1 respectively, such that $b_1(t) = 0$ for almost all $t \leq u_1$, and $b_1(t) = 0$ for almost all $t \geq v_1$. Then

(2.9) is equivalent to the assertion that $u = u_1$ and $v = v_1$, and $b = b_1$ almost everywhere on $[u, v]$. If u_1 and v_1 are defined in this way, and if (2.8) holds, then it follows from Theorem 2.2 that the scalars a, u and v are respectively equal to a_1, u_1 and v_1, and $b_1 = b$ almost everywhere on $\mathcal{I}$.

2.3. Illustration. We conclude this subsection with an example showing that truncated linear models sometimes are, unexpectedly, approximations to rather than equivalent to models that are linear but depend on values taken by X only on a subset of $\mathcal{I}$. In particular, even if g is linear in X on $\mathcal{I}$, and even if g depends on X only through the restriction of X to a subinterval $\mathcal{I}_1$, it may not be possible to represent g as $g(X) = a + \int_{\mathcal{I}_1} b\,X$, for a scalar a and a function b.

Take $\mathcal{I} = [0, 1]$, $u = 0$ and $v = \frac{1}{2}$ for simplicity, and assume that $X(t) = X(1-t)^c$ on $(\frac{1}{2}, 1]$, where $c > 0$ is a constant. Unless $X \geq 0$ on $\mathcal{I}_1 = [0, \frac{1}{2}]$, we should take c to be an integer, but no matter what the sign of c we assume that $c \neq 1$. Let b_1 and b_2 be functions defined on $\mathcal{I}_1$, and put

$$g(X) = \int_{\mathcal{I}_1} \bigl\{ b_1(t)\,X(t) + b_2(t)\,X(t)^c \bigr\}\,dt = \int_{\mathcal{I}} b_3(t)\,X(t)\,dt\,,$$

where $b_3 = b_1$ on $[0, \frac{1}{2}]$ and $b_3(t) = b_2(1-t)$ for $t \in (\frac{1}{2}, 1]$. This formula presents g as a linear model, but one where g depends only on the restriction of X to $\mathcal{I}_1$. Nevertheless it is not, in general, possible to write

$$g(X) = a + \int_{\mathcal{I}_1} b(t)\,X(t)\,dt\,,$$

for a scalar a and function b.

3. METHODOLOGY

3.1. Methodology for estimating a, b and interval endpoints. We suggest methodology in the case $u = 0$, which is the practical setting that motivated our work. In that context we write θ for v. It is assumed that we have independent data pairs (X_i, Y_i), for $1 \leq i \leq n$, all distributed as (X, Y).

Let $\psi_1, \psi_2, \ldots$ denote an orthonormal basis for the class of square-integrable functions on $\mathcal{I}$. Then we can write

$$b(t) = \sum_{j=1}^{\infty} \beta_j\,\psi_j(t)\,, \quad X(t) = \sum_{j=1}^{\infty} \xi_j\,\psi_j(t)\,, \quad \int_{\mathcal{I}} b\,X = \sum_{j=1}^{\infty} \beta_j\,\xi_j\,, \tag{3.1}$$

where $\beta_j = \int_{\mathcal{I}} b\,\psi_j$, $\xi_j = \int_{\mathcal{I}} X\,\psi_j$, and the first two series in (3.1) represent generalised Fourier representations for $b(t)$ and $X(t)$, respectively. If we truncate the third series after m terms then we obtain an approximation to the regression mean:

$$E(Y \,|\, X) \approx a + \sum_{j=1}^{m} \beta_j\,\xi_j\,.$$

Inference in the linear functional regression model often is based on this generalised Fourier approximation. In practice the ψ_js are often chosen to be empirical principal component functions, for example the functions $\hat\phi_j$ defined by the singular-value decomposition of the empirical covariance function:

$$\widehat K(t_1,t_2) = \frac{1}{n}\sum_{i=1}^{n}\{X_i(t_1)-\bar X(t_1)\}\,\{X_i(t_2)-\bar X(t_2)\} = \sum_{j=1}^{\infty}\widehat\omega_j\,\hat\phi_j(t_1)\,\hat\phi_j(t_2)\,, \quad (3.2)$$

where $(\widehat\omega_j, \hat\phi_j)$ are the (eigenvalue, eigenfunction) pairs associated with the linear operator with kernel $\widehat K$, $\bar X = n^{-1}\sum_i X_i$, and terms are ordered such that $\widehat\omega_1 \ge \widehat\omega_2 \ge \ldots$.

The expansion (3.2) is an empirical version of (2.7), and $\widehat\omega_j$ and $\hat\phi_j$ are, under mild conditions, root-n consistent estimators of ω_j and ϕ_j, respectively, in (2.7). (See Hall and Hosseini-Nasab, 2009.) If the random functions X_i are continuous, in the sense that the expected value of the Lebesgue measure set $\{t \in \mathcal{I} : X_1(t) = X_2(t)\}$ equals 0, then with probability 1 the functions $\hat\phi_1, \ldots, \hat\phi_n$ are orthonormal on $\mathcal{I}$, reflecting the fact that $\phi_1, \phi_2, \ldots$ in (2.7) are orthonormal on $\mathcal{I}$. Since $\widehat\omega_j = 0$ for $j \ge n+1$ then the $\hat\phi_j$s are not defined explicitly for j in this range.

In the case of the truncated linear functional regression model, at least two approaches are feasible, as follows.

A. First method: Simultaneous inference. Determine estimators $\hat a$, $\hat\beta_j$ and $\hat\theta$ of a, β_j and θ by minimising the sum of squares,

$$S_1(a,\beta_1,\ldots,\beta_m,\theta\,|\,m) = \sum_{i=1}^{n}\left[Y_i - a - \int_0^{\theta}\left\{\sum_{j=1}^{m}\beta_j\,\psi_j(t)\right\}X_i(t)\,dt\right]^2. \quad (3.3)$$

Here m can be viewed as a smoothing, or regularisation, parameter for estimating b; taking m too large produces an estimator, $\hat b = \sum_{1\le j\le m}\hat\beta_j\,\psi_j$, that suffers from excessive variance, while choosing m too small results in unnecessarily large bias. A truncated linear predictor of Y, when $X = x$, is given by

$$\hat y(x) = \hat a + \int_0^{\hat\theta}\hat b(t)\,x(t)\,dt\,,$$

being an estimator of $y(x) = a + \int_{[0,\theta]} b\, x$.

Not unexpectedly, however, this approach is inadequate for estimating θ, since it does not encourage the choice of an estimator $\hat\theta$ that is noticeably less than the upper endpoint of the interval $\mathcal{I}$. To improve performance in this regard we add a penalty term to S, obtaining:

$$S(a, \beta_1, \ldots, \beta_m, \theta \,|\, m, \lambda) = S_1(a, \beta_1, \ldots, \beta_m, \theta \,|\, m) + n\,\lambda\,\theta^2\,, \tag{3.4}$$

where $\lambda > 0$ is another tuning parameter. We have multiplied λ by n in (3.4) since both S and S_1 are of order n. The multiplier will assist our intuition when we assess the impact of λ, particularly in section 6. The use of the penalty θ^2 in (3.4) is motivated by a Laplace approximation employed in the methods for selecting λ in Section 3.3, but any continuous increasing function of $\tilde a$ could be used.

If λ is too large then minimising S tends to produce a relatively small estimator $\hat\theta$, whereas if λ is too small then we produce results similar to those obtained by minimising S_1, rather than S; that is, $\hat\theta$ is too large. Choice of m and λ is discussed in section 3.3.

B. Second method: Iterative inference. Here we suggest estimating a and b first, obtaining $\breve a$ and $\breve b$, say, constructed using a standard method; and then estimating θ. Approaches that can be used to compute $\breve a$ and $\breve b$ are discussed by, for example, Ramsay and Silverman (2005, Chapter 12), Hall and Hosseini-Nasab (2006) and Crambes et al. (2008, 2009); see section 1 for a more detailed account of the literature. In the second step for this method we employ again penalised least-squares, but this time we select $\theta = \hat\theta$ to minimise

$$T(\theta) = \sum_{i=1}^n \left\{ Y_i - \breve a - \int_0^\theta \breve b(t)\, X_i(t)\, dt \right\}^2 + n\,\lambda\,\theta^2\,. \tag{3.5}$$

See section 3.3 for choice of λ.

Having computed $\hat\theta$ we can proceed in at least two ways. Most simply, assuming for notational clarity that $\mathcal{I} = [0,1]$, we can define $\hat b$ by truncation and $\hat a$ by correcting $\breve a$ in the obvious way for location:

$$\hat b(t) = \begin{cases} \breve b(t) & \text{if } t \le \hat\theta \\ 0 & \text{if } t > \hat\theta\,, \end{cases} \qquad \hat a = \breve a + \int_0^{\hat\theta} \hat b\, \bar X - \int_0^1 \breve b\, \bar X = \breve a - \int_{\hat\theta}^1 \hat b\, \bar X\,.$$

Alternatively we can use a standard method (for example, the one that produced $\breve a$ and $\breve b$ in (3.5)) to compute new estimators of a and b, this time in the linear regression model $E(Y \,|\, X) = a + \int_0^{\hat\theta} b\, X$.

3.2. Examples of standard methods for estimating a and b. Given a complete orthonormal sequence $\psi_1, \psi_2, \ldots$, and an integer $m \geq 1$, the scalar a and function b typically are defined by minimising

$$\sum_{i=1}^{n} \left[Y_i - a - \int_{\mathcal{I}} \left\{ \sum_{j=1}^{m} \beta_j \, \psi_j(t) \right\} X_i(t) \, dt \right]^2 ;$$

compare (3.3). This results in $\check{b} = \sum_{j \leq m} \check{\beta}_j \, \psi_j$ and $\check{a} = \bar{Y} - \int_{\mathcal{I}} \check{b} \, \bar{X}$, where $\bar{X}$ is as defined in section 3.1, $\bar{Y} = n^{-1} \sum_i Y_i$, and $\check{\beta}_1, \ldots, \check{\beta}_m$ solve the linear system of equations

$$\sum_{j=1}^{m} \beta_j \int_{\mathcal{I}} \int_{\mathcal{I}} \psi_k(t_2) \, \psi_j(t_1) \, \widehat{K}(t_1, t_2) \, dt_1 \, dt_2 = \int_{\mathcal{I}} R(t) \, \psi_k(t) \, dt \tag{3.6}$$

for $k = 1, \ldots, m$, with

$$R(t) = \frac{1}{n} \sum_{i=1}^{n} (Y_i - \bar{Y}) \, \{ X_i(t) - \bar{X}(t) \} .$$

Typically m is determined by cross-validation or an information criterion.

If we take $\psi_j = \hat{\phi}_j$ for $j = 1, \ldots, m$, where $m \leq n$ and $\hat{\phi}_j$ is as in (3.2), then in view of the orthonormality of those functions, (3.6) simplifies conveniently to

$$\beta_k = \check{\beta}_k = \frac{1}{\widehat{\omega}_k} \int_{\mathcal{I}} R(t) \, \hat{\phi}_k(t) \, dt \, ,$$

where $\widehat{\omega}_k$ is as in (3.2). Equivalently,

$$\check{\beta}_k = \frac{1}{n \, \widehat{\omega}_k} \sum_{i=1}^{n} \{ E(Y_i \mid X_i) + \epsilon_i \} \int_{\mathcal{I}} (X_i - \bar{X}) \, \hat{\phi}_k = B_k + \frac{1}{n} \sum_{i=1}^{n} B_{ik} \, \epsilon_i \, , \tag{3.7}$$

where $B_k = (n \, \widehat{\omega}_k)^{-1} \sum_i E(Y_i \mid X_i) \int_{\mathcal{I}} (X_i - \bar{X}) \, \hat{\phi}_k$ and $B_{ik} = \widehat{\omega}_k^{-1} \int_{\mathcal{I}} (X_i - \bar{X}) \, \hat{\phi}_k$, and where $B_k, B_k, \ldots, B_{kn}$ are all measurable in the sigma-field, $\mathcal{X}$ say, generated by $X_1, \ldots, X_n$.

The latter technical property has helpful implications, both practical and theoretical. In regression we undertake inference conditional on the design variables, and so the only source of variability comes from the experimental errors ϵ_i. Result (3.7) tells us that the estimated Fourier components depend linearly in the ϵ_is, with coefficients depending only on the X_is, and in particular that $E(\check{\beta}_j \mid \mathcal{X}) = B_j$ and

$E(\check{b}\,|\,\mathcal{X}) = \sum_{j\le m} B_j\,\hat{\phi}_j$. This ensures a simple, equivalent reformulation of $T(\theta)$, at (3.5):

$$\begin{aligned}
T(\theta) &= \sum_{i=1}^n \left\{ E(Y_i\,|\,X_i) - \check{a} - \int_0^\theta \check{b}(t)\,X_i(t)\,dt \right\}^2 \\
&\quad + 2\sum_{i=1}^n \left\{ E(Y_i\,|\,X_i) - \check{a} - \int_0^\theta \check{b}(t)\,X_i(t)\,dt \right\} (\epsilon_i - \bar{\epsilon}) + \sum_{i=1}^n (\epsilon_i - \bar{\epsilon})^2 + n\,\lambda\,\theta^2 \\
&= \sum_{i=1}^n \left\{ E(Y_i\,|\,X_i) - \check{a} - \int_0^\theta \check{b}(t)\,X_i(t)\,dt \right\}^2 \\
&\quad + 2\sum_{i=1}^n \left\{ E(Y_i\,|\,X_i) - \sum_{j=1}^m B_j \int_0^\theta \hat{\phi}_j(t)\,X_i(t)\,dt \right\} (\epsilon_i - \bar{\epsilon}) \\
&\quad - \frac{2}{n}\sum_{i=1}^n \epsilon_i\,(\epsilon_i - \bar{\epsilon}) \sum_{j=1}^m B_{ij} \int_0^\theta \phi_j(t)\,X_i(t)\,dt + \sum_{i=1}^n (\epsilon_i - \bar{\epsilon})^2 + n\,\lambda\,\theta^2\,. \quad (3.8)
\end{aligned}$$

3.3. Algorithms for selecting tuning parameters. The tuning parameter λ plays an important role regularizing the choice of θ. However, because we expect that $\check{b}(t)$ will be near zero when $t > \theta$, standard methods for choosing λ are unlikely to yield good performance results.

Instead, we propose selecting λ based on our ability to reconstruct a parametric model $\tilde{b}_{\rm simp}(t)$ intended to approximate $b(t)$. We begin by computing an approximate mean squared error for reconstructing $\tilde{b}_{\rm simp}(t)$, using each of Method A and Method B in Section 3.1 for each λ. We then choose the λ that minimises this error and apply it in the original problem. As we shall show in Section 7, our choice of the parametric form for $\tilde{b}_{\rm simp}(t)$ has little effect on the resulting estimators.

After selecting λ, the number of orthogonal components, m, to be employed when using Method A is selected by BIC. In the case of Method B this selection takes place within the estimate of $\check{b}(t)$.

To be specific, implementation of our method for Method A involves the following steps. Method B differs only in Steps (A4) through (A6):

(A1) Construct the pilot estimators $\check{a}$ and $\check{b}(t)$ without a truncation constraint—these will later be employed in Method B.

(A2) From the estimators $\check{a}$ and $\check{b}(t)$, compute residuals $\tilde{\epsilon}_i = Y_i - \check{a} - \int_{\mathcal{I}} \check{b} X_i$ and compute an empirical variance $\hat{\sigma}^2 = n^{-1}\,\sum_{i=1}^n \tilde{\epsilon}_i^2$.

(A3) Compute a parametric estimate $\tilde{b}_{\rm simp}$. This can be a straight line or parametric

curve designed "by eye" to mimic $\check{b}(t)$ and to decrease to 0 and strike the t axis at $\bar{\theta}$, say.

In our simulation analysis in section 4 we shall use a parametric approximation in terms of a Fourier basis:

$$\tilde{b}_{\text{simp}}(t) = \left\{c_0 + c_1 \sin\left(2^k \pi t\right) + c_2 \cos\left(2^k \pi t\right)\right\} I(t < \theta)\,,$$

where $k = 1$, 2 or 3 and we estimate (c_0, c_1, c_2, θ) by minimising unpenalised squared error; see (3.3). We could also choose other low-dimensional representations, such as a polynomial basis. For a fixed-dimensional representation, it is easy to show that these estimates are asymptotically unbiased if the true b falls within the model class.

(A4) For each θ we can obtain expressions for the mean squared prediction error and mean squared error for $\hat{b}$ at truncation point θ when using data from the parametric model,

$$Y_i^* = \check{a} + \int_{\mathcal{I}} \tilde{b}_{\text{simp}}\, X_i + \epsilon_i^*\,, \tag{3.9}$$

where the ϵ_i^*s are distributed as normal $\mathrm{N}(0, \hat{\sigma}^2)$. We shall also use the noiseless expected values $\bar{Y}_i^* = \check{a} + \int_{\mathcal{I}} \tilde{b}_{\text{simp}}\, X_i$.

Specifically, we employ $\psi_j = \hat{\phi}_j^{\theta}$, the empirical Fourier components for the truncated functions $X_i(t)\, I(t < \theta)$ (i.e. the X_is restricted to the range $[0, \theta]$) which have associated variance components τ_j^{θ}. We then obtain estimates a_θ^* and b_θ^* by minimising the squared error for predicting the noiseless data $\bar{Y}_i^*$, and define predicted values $\hat{Y}_i^*(\theta) = a_\theta^* + \int_{\mathcal{I}} b_\theta^*\, X_i$.

(A5) For $\hat{a}_\theta^*$ and $\hat{b}_\theta^*$ estimated from the (hypothetical) data Y_i^*, the mean squared prediction error is

$$S_Y^A(\theta) = E\left\{ \sum_{i=1}^{n} \left(Y_i^* - \hat{a}_\theta^* - \int_{\mathcal{I}} \hat{b}_\theta^* X_i \right)^2 \;\middle|\; \mathcal{X} \right\} = \sum_{i=1}^{n} \left\{ \bar{Y}_i - \hat{Y}_i^*(\theta) \right\}^2 + \hat{\sigma}^2\,(m+1)\,,$$

and the mean squared error for estimating $\tilde{b}_{\text{simp}}$ is given by

$$S_b^A(\theta) = E\left\{ \int_{\mathcal{I}} \left(\tilde{b}_{\text{simp}} - \hat{b}_\theta^* \right)^2 \;\middle|\; \mathcal{X} \right\} = \int_{\mathcal{I}} \left(\tilde{b}_{\text{simp}} - b_\theta^* \right)^2 + \hat{\sigma}^2 \sum_{j=1}^{m} \left(\tau_j^{\theta} \right)^{-1}.$$

(A6) For every λ we select θ to minimise the expected value of (3.4):

$$\theta_\lambda = \text{argmin}_\theta \left\{ S_Y^A(\theta) + \lambda\, \theta^2 \right\}.$$

Further, estimate the variance of this choice by

$$V(\lambda) = \frac{\hat{\sigma}^2}{n} \, \frac{(d^2/d\theta^2)\, S_Y^A(\theta_\lambda)}{\left\{(d^2/d\theta^2)\, S_Y^A(\theta_\lambda) + \lambda\right\}^2} \, .$$

(A7) We now choose λ to minimise the expectation of $S_b^A(\theta)$ with respect to a normal distribution for θ with mean θ_λ and variance $V(\lambda)$:

$$P_b(\lambda) = \int_{\mathcal{I}} S_b^A(\theta) \, \frac{1}{\sqrt{2\pi V(\lambda)}} \, \exp\left\{ -(\theta - \theta_\lambda)^2 / 2V(\lambda) \right\} d\theta \, .$$

For this λ we determine $\hat{a}$ and $\hat{b}$ from the original data via Method A for each choice of m. We repeat the process above for each m, and select m by minimising BIC:

$$\mathrm{BIC}(m) = \log\left\{ \frac{1}{n} \sum_{i=1}^{m} \left(Y_i - \hat{a} - \int_{\mathcal{I}} \hat{b}\, X_i \right)^2 \right\} + (m+1)\, \log(n) \, .$$

For Method B, the process is analogous. We select m via BIC when computing estimates $\check{a}$ and $\check{b}$ in Step (A1). Steps (A4) through (A6) are replaced by:

(B4) Obtain estimates $\check{a}^*$ and $\check{b}^*$ to estimate the $\bar{Y}_i^*$ without truncation using $\psi_j = \hat{\phi}_j^1$—the empirical Fourier components on the full interval with m chosen as in Step (A1), and set the predicted values at θ to be $\check{Y}_i^*(\theta) = \check{a}^* + \int_0^\theta \check{b}^* X_i$.

(B5) Obtain the quantities analogous to $S_Y^A(\theta)$ and $S_b^A(\theta)$ by employing Method B:

$$S_Y^B(\theta) = \sum_{i=1}^{n} \left\{ \bar{Y}_i - \check{Y}_i^*(\theta) \right\}^2 + \hat{\sigma}^2 \sum_{j=1}^{m} \left(\tau_j^1\right)^{-1} \sum_{i=1}^{n} \left(\int_0^\theta \phi_j^1 X_i \right)^2$$

and

$$S_b^B(\theta) = \int_{\mathcal{I}} \left\{ \tilde{b}_{\mathrm{simp}} - \check{b}^* I(t < \theta) \right\}^2 + \hat{\sigma}^2 \sum_{j=1}^{m} \left(\tau_j^\theta\right)^{-1} \int_0^\theta \left(\hat{\phi}_j^1 \right)^2 .$$

(B6) For every λ we select θ to be the first minimum of the expected value of (3.4):

$$\theta_\lambda = \text{first minimum of } S_Y^B(\theta) + \lambda\theta^2 \, ,$$

and estimate the variance of this choice by

$$V(\lambda) = \frac{\hat{\sigma}^2}{n} \, \frac{(d^2/d\theta^2)\, S_Y^B(\theta_\lambda)}{\left\{(d^2/d\theta^2)\, S_Y^B(\theta_\lambda) + \lambda\right\}^2} \, .$$

We suggest using the first minimum of $S_Y^B(\theta)$ because $\check{b}$ is already estimated by minimising squared error. There tends to be a sharp drop in $S_Y^B(\theta)$ close to the

right hand endpoint of the interval, so that $S_Y^B(\theta)+\lambda\theta^2$ is minimised at $\theta = 1$ unless λ is very large.

This scheme is intended to mimic simulating from the model at (3.9), either directly or by bootstrapping the $\tilde{\epsilon}_i$s, but it substantially reduces computational effort. Empirically, $P_b(\lambda)$ approximates the mean squared error for estimating b after choosing the truncation level very well.

Note that we have selected λ based only on our ability to estimate $\tilde{b}_{\text{simp}}$. We could also have included mean squared error for a, and approximated the error for estimating θ within these calculations as well. Observe too that, while these estimation schemes appear similar, the need for a separate principal components analysis for each θ in Method A represents a significant additional computational cost.

While the methods above rely on empirical orthogonal components calculated from the X_is, other finite-dimensional linear representations for b (e.g. explicitly using a polynomial or trigonometric basis) can be employed with some changes to the form of the expected mean squared error calculations above. However, employing Method A we have found that unless the basis is adapted to each θ, the estimate $\hat{b}$ can become numerically unstable as basis functions designed for the interval $[0, 1]$ become close to collinear when restricted to $[0, \theta]$.

4. SIMULATION STUDIES

We expect that the performance of our methods will be affected strongly by the way in which b tends to zero—a discontinuous drop to zero should be easier to detect, while smooth convergence will make it harder to localise the value of θ. To examine the performance of our methods, we conducted a simulation study. In this study we generated covariates X_i, $i = 1, \ldots, 100$, via a trigonometric basis expansion on the range [0,1] given by:

$$\eta_1 = 1,\ \eta_{2k} = \sin\left(2^k \pi t\right), \quad \eta_{2k+1} \cos\left(2^k \pi t\right).$$

The X_is were generated via linear combinations of the first 25 such functions, with coefficients chosen as independent mean-zero Gaussian random variables with the coefficient for η_k having variance $\exp\{-(k-1)/4\}$. This produces exponentially-decaying variance components.

Our estimates of $\check{b}$ were computed using the empirical principal components of the generated $X_i(t)$, and we used two through nine of these, choosing the number by BIC. In Method A, we employed a separate principal components analysis for each value of θ.

The methods also rely on a parametric approximation to $\tilde{b}_{\text{simp}}$. For this, we employed the first two components of the trigonometric basis,

$$\tilde{b}_{\text{simp}}(t, \beta_1, \beta_2, \theta) = \{\beta_1\psi_1(t) + \beta_2\psi_2(t)\}\ I_{t\leq\theta}\,.$$

Finally, we considered three simulation settings, each with $\theta = 0.5$—the midpoint of the interval. A true b for each was given by the following parametric models:

Model 1 $b(t) = \psi_1(t)\, I_{t\leq\theta} = I_{t\leq 0.5}$,

Model 2 $b(t) = \psi_2(t)\, I_{t\leq\theta} = \sin(2\pi t) I_{t<0.5}$,

Model 3 $b(t) = \{\psi_3(t) + \psi_1(t)\}\, I_{t\leq\theta} = \{\cos(2\pi t) + 1\}\, I_{t<0.5}$.

These three functions have discontinuities at θ in the 0th, 1st or 2nd derivatives, representing important behaviour at the $t = \theta$ boundary. Each of these was scaled to give a 26.25/1 signal to noise ratio for additive Gaussian noise with unit variance added to the observations.

Note that in this case, for Model 1 and Model 2, the form of $\tilde{b}_{\text{simp}}$ includes the model class, but it does not for Model 3. We also experimented with specifying $\tilde{b}_{\text{simp}}$ using one and three components of the trigonometric basis system, and found that for all three models, the choice of θ, and hence the estimate of $\hat{b}$, was almost always identical for any choice of $\tilde{b}_{\text{simp}}$ on a given data set—although the selection of λ varied somewhat. We have therefore only presented one of these.

We ran 400 simulations for each of the three models using $n = 100$. Table 1 presents the mean and standard deviation of the estimates of θ for Method A and Method B for each model. Here we see that Method B produces estimates that are less variable than those of Method A, but exhibits more bias. The expected degradation of our estimates when b tends to zero more smoothly is apparent in the observed bias towards more truncation, rather than in the variance. Employing a tapered estimate may improve this.

TABLE 1

	Method A		Method B	
	Mean	Std. Dev.	Mean	Std. Dev
Model 1	0.5197	0.0686	0.5094	0.0156
Model 2	0.4933	0.0823	0.4799	0.031
Model 3	0.3734	0.0649	0.3527	0.0185

Mean and standard deviation of estimates of θ from Method A and Method B based on 400 simulations for each of three models in which $b(t) = 0$ for $t > 0.5$.

Turning to the estimate of b itself, Table 2 presents mean integrated squared error and median integrated squared error for both Method A and Method B as well as for $\check{b}$ employing no truncation for all three simulation models. Here we see that employing truncation results in an improvement for both methods, and that Method A improves on Method B for Models 1 and 3, largely due to the bias that $\check{b}$ exhibits when trying to estimate a function that is identically zero on part of its domain. There is an important distinction in performance for Method A between mean and median squared error, particularly in Model 3; this is due to BIC occasionally selecting a very large m, yielding a high-variance estimate that can be diagnosed readily.

TABLE 2

Mean Squared Error

	Method A	Method B	No Trunc.
Model 1	572.6798	800.2337	1153.9327
Model 2	548.0986	172.1635	317.7989
Model 3	1139.5829	1037.504	1671.0593

Median Squared Error

	Method A	Method B	No Trunc.
Model 1	441.7294	839.5315	1305.7375
Model 2	154.0908	146.0178	275.448
Model 3	495.4372	1131.5556	1696.8166

Mean integrated squared error (left hand columns) and median integrated squared error (right hand columns) for estimating b for each of three models based on 400 simulations.

To illustrate these results more concretely, Figure 1 presents plots of the results of the simulation for Model 2. We have provided histograms of the estimated $\hat{\theta}$ for

each method, as well as plots of the estimate $\hat{b}$ along with the untruncated $\check{b}$. We see that Method B's difficulties are largely associated with bias towards small values of θ, and that there are occasional "wild" estimates from Method A. These results suggest that Method B can be employed as a computationally inexpensive means of deciding whether truncation should be attempted before going to the expense of employing Method A to provide a new estimate. Results for Models 1 and 3 (not shown) are similar.

FIGURE 1

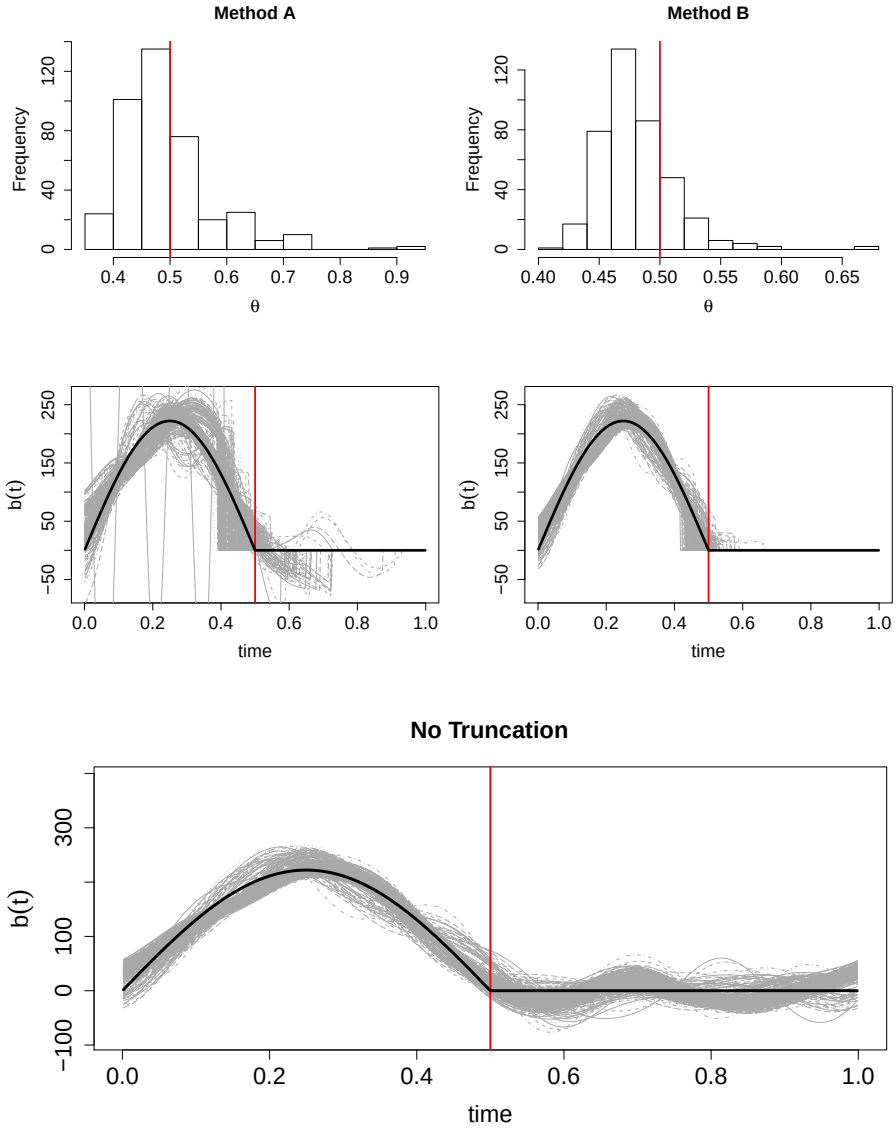

Results of a simulation study with Model 2. Top row: Histograms of $\hat{\theta}$ over 400 simulations; the vertical bar indicates the true values of $\theta = 0.5$. Left column: Method A; Right column: Method B. Second and bottom rows:

estimate $\hat{b}$ (grey lines) along with target (dark, thick lines) for Method A (right), Method B (left) and without truncation (bottom). The vertical bar indicates $\theta = 0.5$.

5. AN ANALYSIS OF PARTICULATE MATTER EMISSIONS

The methods developed in this paper are motivated by a problem of modelling particulate matter (PM) emissions from diesel trucks. For details, see Clark *et. al.* (2007). In these data, trucks are placed on stationary rollers and a particle counter is attached to the exhaust pipe of each. The trucks are then driven through a pre-set driving cycle and PM at the tail pipe is measured every second. Asencio *et. al.* (2014) proposed the following model for these data:

$$\log\{\mathrm{PM}(t)\} = \int_0^\theta b(u)\, Z(t-u)\, du + \epsilon(t)\,,$$

where $Z(t)$ is the acceleration applied by the engine. That is, log(PM) follows a linear model based on the past θ seconds of acceleration. The model is intended to represent mixing of particles in the exhaust pipe. McLean *et. al.* (2014) examined these data for non-linear dependence.

In order to remove dependencies in the data, we have down-sampled PM to obtain an observation every 10 seconds after the first two minutes of data, and have used the previous 60 seconds of acceleration as the corresponding functional covariate. Based on domain knowledge, PM is not expected to take longer than one minute in transport through the exhaust. That is, we have a data set

$$Y_i = \log\{\mathrm{PM}(10i + 120)\}\,,\ X_i(t) = Z(10i + 120 - t)$$

(note that "time" for the stochastic process X_i is now reversed relative to that for Z), where we have 107 observations. The covariates X_i are obtained by smoothing measured velocities in each time window and obtaining a derivative.

Below we illustrate the result of using both Methods A and B to obtain estimates of b and θ. We also perform a residual bootstrap based on the results of Method A, which we used to estimate pointwise standard deviations for each of our estimates and which are represented in the confidence intervals in Figure 2. Here we see that without truncation, the estimate appears to be zero after about 20 seconds. Method B suggests truncating at 18 to 20 seconds, while Method A suggests a possibly longer window although the point estimate for truncation is at 13 seconds. The

roughness of the confidence intervals in Method A are associated with the choice of the number of orthogonal components (re-obtained for each bootstrap) which we chose to be between 2 and 9 via BIC. These results compare with the estimates obtained in Asencio *et. al.* (2014) in which truncation at 40 seconds was selected by cross validation, but where we believe the use of an explicit smoothing penalty may have biassed the results.

FIGURE 2

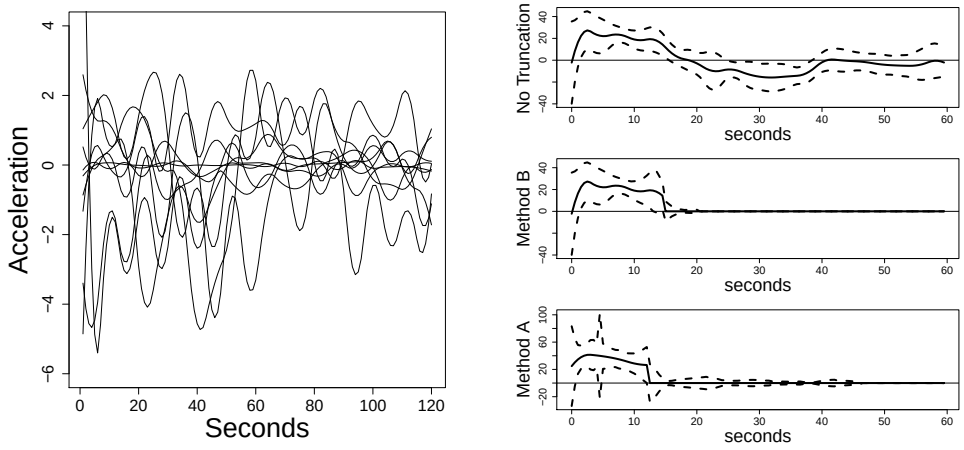

Results of an analysis of PM data. Left: a sample of 10 covariate curves employed in the model. Right from top to bottom: estimates for b with no truncation, using Method B and using Method A. Dashed lines give two-standard-deviation pointwise confidence intervals obtained from 200 residual bootstrap replicates based on the model from Method A.

6. THEORETICAL PROPERTIES

6.1. Main result. In this section we show that the second method suggested in section 3.1 gives consistent estimators of θ_0, denoting the true value of θ. Taking $\mathcal{I} = [0, 1]$, $u = 0$ and $v = \theta_0 \in (0, 1]$, we assume that the truncated linear model at (1.2) is correct. We write a_0 and b_0 for the true values of the scalar a and function b, and assume of b_0 that:

$$b_0 \text{ is continuous on } \mathcal{I}, \text{ vanishes on } [\theta_0, 1] \text{ and is zero at at most a finite number of points in } [0, \theta_0]. \tag{6.1}$$

Taking θ_1, θ_2 to satisfy $0 \leq \theta_1 < \theta_0 < \theta_2 \leq 1$, we further assume that

$$\sup_{\theta_1 \leq \theta \leq \theta_2} \frac{1}{n} \sum_{i=1}^{n} \left[\int_0^{\theta} \left\{ \check{b}(t) - b_0(t) \right\} \left\{ X_i(t) - \bar{X}(t) \right\} dt \right]^2 = O_p\left(\eta_n^2\right), \tag{6.2}$$

where the positive, deterministic sequence η_n satisfies $\eta_n = o(1)$ and $n^{-1} = O(\eta_n^2)$ as $n \to \infty$. The first of these conditions on η_n merely reflects the consistency of $\check{b}$ for b_0, and the second is particularly mild since we would not expect η_n to converge to zero faster than $n^{-1/2}$.

On occasion we suppose in addition that

$$\sup_{\theta_1 \le \theta \le \theta_2} \left| \frac{1}{n} \sum_{i=1}^{n} (\epsilon_i - \bar{\epsilon}) \int_0^\theta \left\{ \check{b}(t) - b_0(t) \right\} \left\{ X_i(t) - \bar{X}(t) \right\} dt \right| = O_p\big(n^{-1/2}\, \eta_n\big)\,. \quad (6.3)$$

In the case of the standard methods discussed in section 3.2, arguments similar to those of Hall and Hosseini-Nasab (2009) can be used to prove that (6.2) and (6.3) hold for the same η_n.

The linear model is assumed to be fitted by least-squares, and so the intercept, a, is estimated by $\check{a} = \bar{Y} - \int_{\mathcal{I}} \check{b}\, \bar{X}$. In particular, $\tilde{a}$ is determined once we have computed the estimator $\tilde{b}$, and the expected value of Y_i, conditional on X_i, is modelled as $\bar{Y} + \int_{\mathcal{I}} \check{b}\, (X_i - \bar{X})$. Therefore it is not necessary to impose analogues of the constraints (6.2) and (6.3) on $\check{a}$ as well as on $\check{b}$.

We suppose too that the errors ϵ_i are independent and identically distributed, and are independent of the X_is; and that second moments are finite:

$$\int_{\mathcal{I}} E\big(X^2\big) + E\big(\epsilon^2\big) < \infty\,, \quad E(\epsilon) = 0\,. \quad (6.4)$$

Then the singular-value decomposition at (2.7) is well defined, with eigenvalues and eigenvectors ω_j and ϕ_j, respectively. We assume that the eigenvalues, which form a nonincreasing sequence of positive numbers, decay sufficiently fast to ensure that

$$\sum_{j=1}^{\infty} \omega_j^{1/2} \sup_{t \in \mathcal{I}} |\phi_j(t)| < \infty\,. \quad (6.5)$$

Let $\hat{\theta}$ denote any value of θ that minimises $T(\theta)$, at (3.5), on the interval $[\theta_1, \theta_2]$, where θ_1 and θ_2 are as in (6.2) and (6.3).

Theorem 6.1. *(i) If* (2.8), (6.1), (6.2), (6.4) *and* (6.5) *hold, and if the penalty parameter* $\lambda = \lambda(n)$ *in* (3.5) *satisfies* $\lambda \to 0$ *and* $\lambda/\eta_n \to \infty$ *as* $n \to \infty$, *then* $\hat{\theta}$ *converges in probability to* θ_0 *as* $n \to \infty$. *(ii) If* (2.8) *and* (6.1)–(6.5) *obtain, and if the penalty parameter* λ *in* (3.5) *satisfies* $\lambda \to 0$ *and* $\lambda/\eta_n^2 \to \infty$ *as* $n \to \infty$, *then the conclusion of part (i) again holds.*

6.2. Discussion. The reason for stating Theorem 6.1 in two versions, one when (6.3) holds and the other without imposing that constraint, is that the conditions on λ are less stringent in the presence of (6.3). In particular, λ can decrease to zero more quickly when (6.3) holds, reflecting the fact that the assumption $\lambda/\eta_n \to \infty$ in part (i) of the theorem is relaxed to $\lambda/\eta_n^2 \to \infty$ in part (ii).

Condition (6.2) typically holds for a choice of η_n that decreases at a polynomial rate in n^{-1} as $n \to \infty$. This reflects the fact that, in a large class of problems, prediction under the linear model can be undertaken with a polynomial level of accuracy, even if we do not have parametric models for the function b_0 or for the distributions of the functions X_i or errors ϵ_i. See Cai and Hall (2006), particularly their Theorem 3.1.

The methods used by Cai and Hall (2006) can be employed to establish (6.2) for $\eta_n = n^{\eta-(1/2)}$, for any given $\eta \in (0, \frac{1}{2})$, provided that, for example, the random functions X have sufficiently many finite moments; the eigenvalues ω_j in (2.7) decrease to zero in asymptotic proportion to j^{-c}, where c is sufficiently large (depending on η); the spacings $\omega_j - \omega_{j+1}$ of the eigenvalues are no less than a fixed constant multiple of $j^{-(c+1)}$; the function b_0 admits a sufficiently rapidly convergent generalised Fourier expansion in terms of the eigenfunctions ϕ_j; and the orthonormal functions ψ_j are taken to be the empirical versions, $\hat{\phi}_j$, of the ϕ_js.

When using conventional methods to compute $\check{b}$, as outlined in section 3.2; and under the assumptions discussed in the previous paragraph; both (6.2) and (6.3) can be shown to hold for values η_n that decrease at a polynomial rate, in particular at rate $n^{\delta-(1/2)}$ where $\delta \in (0, \frac{1}{2})$. We outline details in the last paragraph of this section.

It is possible to establish an upper bound to the rate of convergence of $\hat{\theta}$ to θ_0. The upper bound, and also the actual convergence rate, decreases to zero more slowly, as $n \to \infty$, if we decrease the rate of convergence of $b_0(t)$ to zero as $t \uparrow \theta_0$; that is, if b_0 becomes smoother at θ_0. Here, smoothness can be characterised in terms of the number of bounded derivatives enjoyed by b_0 at θ_0; the greater the number, the slower the convergence rate of $\hat{\theta}$ to θ_0.

To appreciate how (6.3) can be proved if $\check{b}$ is constructed as suggested in section 3.2, recall that $\check{b} = \sum_{k \le m} \check{\beta}_k \, \psi_k$, and note the formula for $\check{\beta}_j$ at (3.7). In that

notation, the quantity being bounded at (6.3) is given by

$$\begin{aligned}
&\sum_{i=1}^n (\epsilon_i - \bar\epsilon) \int_0^\theta \{\check b(t) - b_0(t)\}\, \{X_i(t) - \bar X(t)\}\, dt \\
&= \sum_{i=1}^n \epsilon_i \int_0^\theta \left\{ \sum_{k=1}^m \left(B_k + \frac{1}{n} \sum_{i_1=1}^n B_{i_1 k}\, \epsilon_{i_1} \right) \psi_k(t) - b_0(t) \right\} \{X_i(t) - \bar X(t)\}\, dt \\
&= \sum_{i=1}^n \epsilon_i \int_0^\theta \left\{ \sum_{k=1}^m B_k\, \psi_k(t) - b_0(t) \right\} \{X_i(t) - \bar X(t)\}\, dt \\
&\quad + \frac{1}{n} \sum_{i=1}^n \epsilon_i^2 \int_0^\theta \left\{ \sum_{k=1}^m B_{ik}\, \psi_k(t) \right\} \{X_i(t) - \bar X(t)\}\, dt \\
&\quad + \sum_{i=1}^n \epsilon_i \int_0^\theta \left\{ \left(\frac{1}{n} \sum_{i_1 \,:\, i_1 \ne i} B_{i_1 k}\, \epsilon_{i_1} \right) \psi_k(t) \right\} \{X_i(t) - \bar X(t)\}\, dt\,, \qquad (6.6)
\end{aligned}$$

where all but the quantities that are written explicitly as ϵ_i or ϵ_{i_1} are measurable in the sigma-field generated by $\mathcal{X}$, and so are conditioned upon and are independent of the errors $\epsilon_1, \ldots, \epsilon_n$. In consequence, the methods developed by Cai and Hall (2006) can be used to establish (6.3).

A. TECHNICAL ARGUMENTS

A.1. Proof of Theorem 2.1. Since, in view of (1.4), $D_2(h) = D_3(h) + E(\epsilon^2)$, where $D_3(h) = E\{g(X) - h(X)\}^2$, then h is chosen equivalently to minimise $D_3(h)$; and since, by (2.2), h can be translated to an arbitrary extent, then $E\{g(X) - h_0(X)\} = 0$ at the minimum, i.e. the first part of (2.5) must hold. Hence, without loss of generality, $E\{g(X)\} = 0$. Therefore we wish to minimise

$$D_3(h) = E\{g(X) - h(X)\}^2 = E\{g(X)^2\} - 2\, E\{g(X)\, h(X)\} + E\{h(X)^2\}\,,$$

in the class $\mathcal{H}_1 \subseteq \mathcal{H}$ of functions $h \in \mathcal{H}$ such that $E\{h(X)\} = 0$. We claim that any $h_0 \in \mathcal{H}_1$ that minimises $D_2(h_1)$ over that class satisfies the second part of (2.5).

To appreciate why, suppose $h = h_0$ produces a minimum, and put $h_1 = c\, h$ where c is a constant. (Clearly, $h_1 \in \mathcal{H}_1$ whenever $h \in \mathcal{H}_1$.) Then

$$D_3(h_1) = E\{g(X)^2\} - 2\, c\, E\{g(X)\, h(X)\} + c^2\, E\{h(X)^2\}\,.$$

Now, $D_3(h_1)$, treated as a function of c, is a convex parabola, and achieves a unique minimum when

$$c = c_1 \equiv \frac{E\{g(X)\, h(X)\}}{E\{h(X)^2\}}\,.$$

The second part of (2.5) must hold if h there is replaced by $h_2 = c_1\, h$. This contradicts our assumption that h produces a minimum, unless of course $c_1 = 1$, in which case the second part of (2.5) obtains. (That result is equivalent to $c_1 = 1$.)

A.2. Proof of Theorem 2.2. Since any adjustment that centres X can be accommodated in the scalars a and a_1 in (2.6), then we can assume without loss of generality that $E(X) = 0$. Then (2.6) implies that $a = a_1$ and

$$E\Big(\int_{\mathcal{I}} \delta\, X\Big)^2 = 0\,, \tag{A.1}$$

where δ denotes the function defined by $\delta(t) = b(t)\, I(t \in [u, v]) - b_1(t)\, I(t \in [u_1, v_1])$. Now, the left-hand side of (A.1) is given by

$$\int_{\mathcal{I}}\int_{\mathcal{I}} \delta(t_1)\, \delta(t_2)\, K(t_1, t_2)\, dt_1\, dt_2 = \sum_{j=1}^{\infty} \omega_j \left(\int_{\mathcal{I}} \delta\, \phi_j\right)^2,$$

where we used (2.7) to derive the identity. Therefore (A.1) holds if and only if, for each j, $\omega_j^{1/2} \int_{\mathcal{I}} \delta\, \phi_j = 0$, and in view of (2.8) this is equivalent to $\delta = 0$ almost everywhere on $\mathcal{I}$.

A.3. Proof of Theorem 6.1.

Step 1: Preparatory lemma. Let θ_1, θ_2 be as in (6.2) and (6.3). That is, $0 \le \theta_1 < \theta_0 < \theta_2 \le 1$.

Lemma. *If* (6.1), (6.4) *and* (6.5) *hold then, uniformly in* $\theta \in [\theta_1, \theta_2]$,

$$\sum_{i=1}^{n} (\epsilon_i - \bar{\epsilon}) \int_{\theta}^{\theta_0} b_0(t)\, \{X_i(t) - \bar{X}(t)\}\, dt = O_p\big(n^{1/2}\, |\theta - \theta_0|\big)\,. \tag{A.2}$$

Proof. We show that, uniformly in $\theta \in [\theta_1, \theta_2]$,

$$\sum_{i=1}^{n} \epsilon_i \int_{\theta}^{\theta_0} b_0(t)\, \{X_i(t) - \mu(t)\}\, dt = O_p\big(n^{1/2}\, |\theta - \theta_0|\big)\,, \tag{A.3}$$

where $\mu = E(X)$. A similar but simpler argument demonstrates that, uniformly in the same sense,

$$\bar{\epsilon} \sum_{i=1}^{n} \int_{\theta}^{\theta_0} b_0(t)\, \{X_i(t) - \mu(t)\}\, dt = O_p\big(n^{1/2}\, |\theta - \theta_0|\big)\,, \tag{A.4}$$

and together, (A.3) and (A.4) imply that

$$\sum_{i=1}^{n} (\epsilon_i - \bar{\epsilon}) \int_{\theta}^{\theta_0} b_0(t) \{X_i(t) - \mu(t)\} \, dt = O_p\big(n^{1/2} \, |\theta - \theta_0|\big) \, ,$$

which, since $\sum_i (\epsilon_i - \bar{\epsilon}) \int_{\theta}^{\theta_0} b_0 \, (\bar{X} - \mu) = 0$, is equivalent to (A.2).

To derive (A.3), first recall the singular-value decomposition at (2.7), involving the (eigenvalue, eigenfunction) pairs (ω_j, ϕ_j). The same decomposition applies to the random function $\epsilon \, (X - \mu)$, except that the eigenvalues are now $\sigma^2 \, \omega_j$, where $\sigma^2 = E(\epsilon^2)$. (The eigenfunctions are unchanged.) Therefore we can write

$$\epsilon \, \{X(t) - \mu(t)\} = \sum_{j=1}^{\infty} \xi_j \, \phi_j(t) \, ,$$

where the random variables ξ_j have zero means and respective variances $\sigma^2 \, \omega_j$. Hence, writing $\bar{\xi}_j$ for the mean of n random variables all independent and identically distributed as ξ_j, we have:

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^{n} \epsilon_i \int_{\theta}^{\theta_0} b_0(t) \, \{X_i(t) - \mu(t)\} \, dt \right| &= \left| \int_{\theta}^{\theta_0} b_0(t) \left\{ \sum_{j=1}^{\infty} \bar{\xi}_j \, \phi_j(t) \right\} dt \right| \\ &\le |\theta - \theta_0| \left\{ \sup_{t \in \mathcal{I}} |b_0(t)| \right\} \sum_{j=1}^{\infty} |\bar{\xi}_j| \sup_{t \in \mathcal{I}} |\phi_j(t)| \, . \end{aligned} \tag{A.5}$$

Writing $s_j = \sup_{t \in \mathcal{I}} |\phi_j(t)|$, we have:

$$E\bigg(\sum_{j=1}^{\infty} |\bar{\xi}_j| \, s_j \bigg) \le \sum_{j=1}^{\infty} \big(E\bar{\xi}_j^2\big)^{1/2} s_j = n^{-1/2} \, \sigma \sum_{j=1}^{\infty} \omega_j^{1/2} \, s_j = O\big(n^{-1/2}\big) \, , \tag{A.6}$$

where the last identity follows from (6.5). Result (A.3) is a consequence of (A.5) and (A.6). Note that the fact that b_0 is continuous on $\mathcal{I}$, as assumed in (6.1), implies that $\sup_{t \in \mathcal{I}} |b_0(t)| < \infty$.

Step 2: Expansions of $T(\theta)$. The two expansions, first in cases where both (2.7) and (6.3) hold, and secondly where only (6.2) obtains, are given at (A.9) and (A.11), respectively.

Since $\mathcal{I} = [0, 1]$ and $Y_i = E(Y) + \int_{[0,\theta_0]} b_0 \, (X_i - EX) + \epsilon_i$ then $Y_i - \bar{Y} = \int_{[0,\theta_0]} b_0 \, (X_i - \bar{X}) + \epsilon_i - \bar{\epsilon}$, and hence,

$$T(\theta) - n \, \lambda \, \theta = \sum_{i=1}^{n} \left\{ Y_i - \breve{a} - \int_0^{\theta} \breve{b}(t) \, X_i(t) \, dt \right\}^2$$

$$
\begin{aligned}
&= \sum_{i=1}^n \biggl[Y_i - \bar{Y} - \int_0^\theta \check{b}(t) \, \{X_i(t) - \bar{X}(t)\} \, dt \biggr]^2 \\
&= \sum_{i=1}^n \biggl[\int_0^{\theta_0} b_0(t) \, \{X_i(t) - \bar{X}(t)\} \, dt - \int_0^\theta \check{b}(t) \, \{X_i(t) - \bar{X}(t)\} \, dt + \epsilon_i - \bar{\epsilon} \biggr]^2 \\
&= \sum_{i=1}^n \biggl[\int_\theta^{\theta_0} b_0(t) \, \{X_i(t) - \bar{X}(t)\} \, dt \\
&\qquad\qquad - \int_0^\theta \bigl\{ \check{b}(t) - b_0(t) \bigr\} \, \{X_i(t) - \bar{X}(t)\} \, dt + \epsilon_i - \bar{\epsilon} \biggr]^2 \\
&= \sum_{i=1}^n \biggl(\biggl[\int_\theta^{\theta_0} b_0(t) \, \{X_i(t) - \bar{X}(t)\} \, dt \biggr]^2 \\
&\qquad\qquad + \biggl[\int_0^\theta \bigl\{ \check{b}(t) - b_0(t) \bigr\} \, \{X_i(t) - \bar{X}(t)\} \, dt \biggr]^2 \biggr) \\
&\qquad\qquad + 2 \sum_{i=1}^n \biggl[\int_\theta^{\theta_0} b_0(t) \, \{X_i(t) - \bar{X}(t)\} \, dt \\
&\qquad\qquad\qquad - \int_0^\theta \bigl\{ \check{b}(t) - b_0(t) \bigr\} \, \{X_i(t) - \bar{X}(t)\} \, dt \biggr] (\epsilon_i - \bar{\epsilon}) + \sum_{i=1}^n (\epsilon_i - \bar{\epsilon})^2 \\
&= T_4(\theta) + \sum_{i=1}^n (\epsilon_i - \bar{\epsilon})^2 + O_p\Bigl[n^{1/2} \bigl\{ |\theta - \theta_0| \, I(\theta) + \eta_n \bigr\} + n \eta_n^2 \Bigr] , \qquad \text{(A.7)}
\end{aligned}
$$

uniformly in $\theta \in [\theta_1, \theta_2]$, where $I(\theta) = I(\theta < \theta_0)$,

$$
T_4(\theta) = n \int_\theta^{\theta_0} \int_\theta^{\theta_0} b_0(t_1) \, b_0(t_2) \, \widehat{K}(t_1, t_2) \, dt_1 \, dt_2 \, , \qquad \text{(A.8)}
$$

and $\widehat{K}$ is as at (3.2). To obtain the last identity in (A.7) we used formulae (6.2), (6.3) and (A.2).

Using the fact that, as assumed immediately below (T.2), $n^{-1} = O(\eta_n^2)$, we deduce that $n^{1/2} \eta_n = O(n \eta_n^2)$. Therefore (A.7) entails:

$$
n^{-1} \, T(\theta) = n^{-1} \, T_4(\theta) + \frac{1}{n} \sum_{i=1}^n (\epsilon_i - \bar{\epsilon})^2 + \lambda \, \theta^2 + O_p\bigl\{ n^{-1/2} \, |\theta - \theta_0| \, I(\theta) + \eta_n^2 \bigr\} \, . \qquad \text{(A.9)}
$$

Of course, (6.3) is assumed only in part (ii) of Theorem 6.1. In the statement of part (i) of the theorem we impose condition (6.2) but not (6.3), but from (6.2) we can derive the following bounds, uniformly in $\theta \in [\theta_1, \theta_2]$, for the quantity of which the absolute value is taken on the left-hand side of (6.3):

$$
\biggl| \sum_{i=1}^n (\epsilon_i - \bar{\epsilon}) \int_0^\theta \bigl\{ \check{b}(t) - b_0(t) \bigr\} \, \{X_i(t) - \bar{X}(t)\} \, dt \biggr|
$$

$$\leq \left(\sum_{i=1}^{n} (\epsilon_i - \bar{\epsilon})^2 \cdot \sum_{i=1}^{n} \left[\int_0^{\theta} \left\{ \check{b}(t) - b_0(t) \right\} \left\{ X_i(t) - \bar{X}(t) \right\} dt \right]^2 \right)^{1/2}$$
$$= O_p\big\{ (n \cdot n\eta_n^2)^{1/2} \big\} = O_p\big(n\eta_n\big) \,. \tag{A.10}$$

Using (A.10) in place of (6.3), but in all other respects using the argument leading to (A.7), we obtain the following expansion in place of (A.9):

$$n^{-1}\, T(\theta) = n^{-1}\, T_4(\theta) + \frac{1}{n} \sum_{i=1}^{n} (\epsilon_i - \bar{\epsilon})^2 + \lambda\, \theta^2 + O_p\Big\{ n^{-1/2}\, |\theta - \theta_0|\, I(\theta) + \eta_n \Big\} \,, \tag{A.11}$$

uniformly in $\theta \in [\theta_1, \theta_2]$.

Step 3: Completion. Combining (A.9) and (A.11) we deduce that

$$n^{-1}\, T(\theta) = n^{-1}\, T_4(\theta) + n^{-1} \sum_{i=1}^{n} (\epsilon_i - \bar{\epsilon})^2 + \lambda\, \theta^2 + O_p\big\{ n^{-1/2}\, |\theta - \theta_0|\, I(\theta) \big\} + \begin{cases} O_p(\eta_n^2) & \text{if (6.2) and (6.3) hold} \\ O_p(\eta_n) & \text{if only (6.2) holds} \,, \end{cases} \tag{A.12}$$

where the remainders are of the stated sizes uniformly in $\theta \in [\theta_1, \theta_2]$. If $\theta > \theta_0$ then, in view of (6.1) and the definition of $T_4(\theta)$ at (A.8), $T_4(\theta) = 0$, and so (A.12) simplifies to:

$$n^{-1}\, T(\theta) = U + \lambda\, (\theta^2 - \theta_0^2) + \begin{cases} O_p(\eta_n^2) & \text{if (6.2) and (6.3) hold} \\ O_p(\eta_n) & \text{if only (6.2) holds} \,, \end{cases}$$

where the random variable U does not depend on θ, and now the remainders are of the stated sizes uniformly in $\theta \in [\theta_0, \theta_2]$. Hence, since $\lambda/\eta_n^2 \to \infty$ if (6.2) and (6.3) both hold (i.e. if we are in the context of part (i) of Theorem 6.1); and since $\lambda/\eta_n \to \infty$ if only (6.2) is assumed (i.e. if we are in the context of part (i) of the theorem); and if $\hat{\theta}$ is chosen to minimise $T(\theta)$ for $\theta \in [\theta_1, \theta_2]$; then, for each $\delta > 0$, $P(\hat{\theta} > \theta_0 + \delta) \to 0$.

On the other hand, if $\theta < \theta_0$ then (A.12) implies that

$$n^{-1}\, T(\theta) = n^{-1}\, T_4(\theta) + V - \lambda\, (\theta_0^2 - \theta^2) + \begin{cases} O_p\{ n^{-1/2}\, (\theta_0 - \theta) + \eta_n^2 \} & \text{if (6.2) and (6.3) hold} \\ O_p\{ n^{-1/2}\, (\theta_0 - \theta) + \eta_n \} & \text{if only (6.2) holds} \,, \end{cases} \tag{A.13}$$

where the random variable V does not depend on θ. Define

$$\kappa(\theta) = \int_{\theta}^{\theta_0} \int_{\theta}^{\theta_0} b_0(t_1)\, b_0(t_2)\, K(t_1, t_2)\, dt_1\, dt_2 \,.$$

In the next paragraph we show that, for any bounded function f on $\mathcal{I}$,

$$\sup_{u,v\in\mathcal{I}} \left| \int_u^v \int_u^v f(t_1)\, f(t_2) \left\{ \widehat{K}(t_1,t_2) - K(t_1,t_2) \right\} dt_1\, dt_2 \right| \to 0\,, \tag{A.14}$$

where the convergence is in probability. Taking $f = b_0$ we deduce from the definition of $T_4(\theta)$ at (A.8) that

$$\sup_{\theta\,:\,\theta\in[\theta_1,\theta_0]} \left| n^{-1}\, T_4(\theta) - \kappa(\theta) \right| \to 0 \tag{A.15}$$

in probability. It follows from (6.1) that $\kappa(\theta)$ is strictly positive whenever $\theta < \theta_0$. Using this property, (A.13), (A.15) and the fact that $\lambda \to 0$ as $n \to \infty$, we deduce that for each $\delta > 0$, $P(\hat\theta < \theta_0 - \delta) \to 0$. Combining this result with the property that $P(\hat\theta > \theta_0 + \delta) \to 0$, derived in the previous paragraph, we deduce that $\hat\theta \to \theta_0$ in probability, as had to be proved.

To derive (A.14), let LHS denote the left-hand side of (A.14), and note that, since $\mathcal{I} = [0,1]$, then LHS $\le (\sup|f|)^2\, \hat J$, where

$$\hat J \equiv \int_{\mathcal{I}} \int_{\mathcal{I}} \left| \widehat{K}(t_1,t_2) - K(t_1,t_2) \right| dt_1\, dt_2\,.$$

Since $\int_{\mathcal{I}} E(X^2) < \infty$ (see (6.4)) then $E\{|\widehat{K}(t_1,t_2) - K(t_1,t_2)|\} \to 0$ as $n \to \infty$, for each pair $t_1, t_2 \in \mathcal{I}$. Similarly, $E(\hat J) \to 0$ as $n \to \infty$. Hence, $\hat J \to 0$ in probability, implying that LHS $\to 0$ in probability, i.e. (A.14) holds.

REFERENCES

APANASOVICH, T.V. AND GOLDSTEIN, E. (2008). On prediction error in functional linear regression. *Statist. Probab. Lett.* **78**, 1807–1810.

ASENCIO, M., HOOKER, G. AND GAO H.O., (2014), Functional Convolution Models. *Stat. Mod.* **14**, 1-21.

BAÍLLO, A. (2009). A note on functional linear regression. *J. Stat. Comput. Simul.* **79**, 657–669.

CAI, T.T. AND HALL, P. (2006). Prediction in functional linear regression. *Ann. Statist.* **34**, 2159–2179.

CAI, T.T. AND YUAN, M. (2012). Minimax and adaptive prediction for functional linear regression. *J. Amer. Statist. Assoc.* **107**, 1201–1216.

CAI, T.T. AND ZHOU, H.H. (2013). Adaptive functional linear regression. `http://www-stat.wharton.upenn.edu/~tcai/paper/Adaptive-FLR.pdf`.

CARDOT, H., FERRATY, F. AND SARDA, P. (1999). Functional linear model. *Statist. Probab. Lett.* **45**, 11–22.

CARDOT, H., FERRATY, F. AND SARDA, P. (2003). Spline estimators for the functional linear model. *Statist. Sinica* **13**, 571–591.

CARDOT, H., MAS, A. AND SARDA, P. (2007). CLT in functional linear regression models. *Probab. Theory Related Fields* **138**, 325–361.

CLARK, N. N., GAUTAM, M., WAYNE, W. S., LYONS, D. W., THOMPSON, G. AND ZIELINSKA, B. (2007). Heavy-duty chassis dynamometer testing for emissions inventory, air quality modeling, source apportionment and air toxins emissions inventory: E55/59 all phases. *Technical Report E55/59, Coordinating Research Council.*

COMTE, F. AND JOHANNES, J. (2012). Adaptive functional linear regression. *Ann. Statist.* **40**, 2765–2797.

CRAMBES, C., KNEIP, A. AND SARDA, P. (2008). Estimation of the functional linear regression with smoothing splines. In *Functional and Operatorial Statistics*, eds S. Dabo-Niang and F. Ferraty, pp. 117–120. *Contrib. Statist.*, Physica-Verlag/Springer, Heidelberg.

CRAMBES, C., KNEIP, A. AND SARDA, P. (2009). Smoothing splines estimators for functional linear regression. *Ann. Statist.* **37**, 35–72.

FAN, J. AND ZHANG, J.-T. (2000). Two-step estimation of functional linear models with applications to longitudinal data. *J. R. Stat. Soc.* Ser. B **62**, 303–322.

FERRATY, F., GONZÁLEZ-MANTEIGA, W., MARTÍNEZ-CALVO, A. AND VIEU, P. (2012). Presmoothing in functional linear regression. *Statist. Sinica* **22**, 69–94.

HALL, P. AND HOROWITZ, J.L. (2007). Methodology and convergence rates for functional linear regression. *Ann. Statist.* **35**, 70–91.

HALL, P. AND HOSSEINI-NASAB, M. (2006). On properties of functional principal components analysis. *J. R. Stat. Soc.* Ser. B **68**, 109–126.

HALL, P. AND HOSSEINI-NASAB, M. (2009). Theory for high-order bounds in functional principal components analysis. *Math. Proc. Cambridge Philos. Soc.* **146**, 225–256.

HE, G., MÜLLER, H.-G. AND WANG, J.-L. (2010). Functional linear regression via canonical analysis. *Bernoulli* **16**, 705–729.

JAMES, G., WANG, J. AND ZHU, J. (2009). Functional linear regression that's interpretable. *Ann. Statist.* **37**, 2083–2108.

JOHANNES, J. AND SCHENK, R. (2012). Adaptive estimation of linear functionals in functional linear models. *Math. Methods Statist.* **21**, 189–214.

JOHANNES, J. AND SCHENK, R. (2013). On rate optimal local estimation in functional linear model. *Electronic J. Statist.* **7**, 191–216.

LI, Y. AND HSING, T. (2007). On rates of convergence in functional linear regres-

sion. *J. Multivariate Anal.* **98**, 1782–1804.

MARONNA, R.A. AND YOHAI, V.J. (2013). Robust functional linear regression based on splines. *Comput. Statist. Data Anal.* **65**, 46–55.

MAS, A. AND PUMO, B. (2009). Functional linear regression with derivatives. *J. Nonparametr. Statist.* **21**, 19–40.

MCLEAN, M. W., HOOKER, G. AND RUPPERT, D., (2014). Restricted Likelihood Ratio Tests for Linearity in Scalar-on-Function Regression. *Stat. Comp.* in press.

RAMSAY, J.O. AND SILVERMAN, B.W. (2002). *Applied Functional Data Analysis*. Springer, New York.

RAMSAY, J.O. AND SILVERMAN, B.W. (2005). *Functional Data Analysis*, Second Edn. Springer, New York.

YAO, F., MÜLLER, H.-G. AND WANG, J.-L. (2005). Functional linear regression analysis for longitudinal data. *Ann. Statist.* **33**, 2873–2903.

YUAN, M. AND CAI, T.T. (2010). A reproducing kernel Hilbert space approach to functional linear regression. *Ann. Statist.* **38**, 3412–3444.

WU, Y., FAN, J. AND MÜLLER, H.-J. (2010). Varying-coefficient functional linear regression. *Bernoulli* **16**, 730–758.